PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

EX PARTE Andrew W. Wilson

Application for Patent

Filed September 17, 2003

Application No. 10/665,846

FOR:

NO SINGLE POINT OF FAILURE RAID BOX USING SATA DRIVES

APPEAL BRIEF

CERTIFICATE OF MAILING
I hereby certify that this correspondence is being deposited with the United States posta Service as First Class mail to Commissioner for Patents PO Box 1450 Alexandria VA, 22313-1450 on August 27, 2009.

Signed: ______________________
Michael L. Gencarella

MARTINE PENILLA & GENCARELLA, LLP
Attorneys for Applicants

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST

The real party in interest is Adaptec, Inc., the assignee of the present application.

II. RELATED APPEALS AND INTERFERENCES

The Applicants are not aware of any related appeals or interferences.

III. STATUS OF CLAIMS

Claims 1-22 and 30-34 are pending in the subject application. Claims 1-22 and 30-34 have been rejected and are on appeal. Claims 23-29 have been cancelled.

IV. STATUS OF AMENDMENTS

Applicants submitted an amendment on September 17, 2008, in response to a non-Final Office Action mailed on March 17, 2008. This amendment was the last entered amendment. A request for Reconsideration was submitted on February 24, 2009 in response to the Final Office action of December 24, 2008.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The subject invention is directed towards an InfiniBand network architecture where router circuitry communicates data between a host and a target device. The router circuitry includes circuitry for generating an external queue pair (QP) for establishing communication between the router circuitry and the host and circuitry for generating internal queue pairs to establish communication between the router circuitry and a device controller, between the between the device controller and the target device, and between

the router circuitry and the target device. In addition, the embodiments provide for a RAID system that avoids having a single point of failure.

Claim 1 defines a method for storing data *(see Figure 17)*, comprising:

transmitting a storage operation request to one of at least two controllers (*Figure 16 108a and 108b, Figure 17 operation 700)*, the at least two controllers being capable of managing communication with a plurality of targets(*Figure 16 721a and 721p)*;

directing the storage operation request to an operational one of at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802)* when the one of at least two controllers is inoperable through an L4 router *(Figure 16, 162a, 162b)*, the directing including,

generating an external queue pair (QP) *(168) (page 16, lines 1-3)* for establishing communication between the L4 router *(IB L4 Router)* and an external host *(102, 104)* through a reliable connection (RC) session *(lines from 168 to hosts 102 and 104) (page 17, lines 10-12) (page 17, lines 19-21)*, and

generating internal queue pairs *(164, pairs connected to EEC boxes) (page 17, lines 21-22) (page 18, lines 1-2)*, one of the internal queue pairs being coupled with the external queue pair*(168) (page 17, lines 10-12; page 18, lines 1-2)*, the internal queue pairs establishing communications between the at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802)* and an internal target device *(112', 114') (page 16, lines 15-16)*, between the L4 router circuitry *(IB L4 Router)* and the at least two controllers, and between the L4 router circuitry *(IB L4 Router)* and the internal target device *(112', 114') (page 16, lines 15-16)* by using reliable datagram (RD) sessions *(page 18, lines 3-5)*; and

processing the storage operation request with the operational one of the at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802).*

Claim 11 defines a method for storing data *(see Figure 17)*, comprising:

providing a plurality of storage containers, each one of the plurality storage containers having a plurality of storage devices *(See Figure 16, storage devices 721a-p and page 46 lines 12-19)*;

generating a plurality of storage volumes where each one of the storage volumes includes at least one storage device from each one of the plurality storage containers *(See Figure 16, storage devices 721a-p and page 46 lines 12-19)*;

managing each of the plurality of volumes with a corresponding storage device controller *(See Figure 16 108a and 108b and page 46 lines 8-14)*; and

when the corresponding storage device controller is inoperable *(see Figure 17 and page 48 lines 8-10),*

accessing data on the plurality of storage volumes through at least one operable storage device controller that is configured to access the volume managed by the inoperable storage device controller *(see Figure 17 and page 48 lines 8-10)*, the accessing through an L4 router, and including,

generating an external queue pair (QP) *(168) (page 16, lines 1-3)* for establishing communication between the L4 router *(IB L4 Router)* and an external host *(102, 104)* through a reliable connection (RC) session *(lines from 168 to hosts 102 and 104) (page 17, lines 10-12) (page 17, lines 19-21)*, and

generating internal queue pairs *(164, pairs connected to EEC boxes) (page 17, lines 21-22) (page 18, lines 1-2)*, one of the internal queue pairs being coupled with the

external queue pair*(168) (page 17, lines 10-12; page 18, lines 1-2)*, the internal queue pairs establishing communications between the at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802)* and an internal target device *(112', 114') (page 16, lines 15-16)*, between the L4 router circuitry *(IB L4 Router)* and the at least two controllers, and between the L4 router circuitry *(IB L4 Router)* and the internal target device *(112', 114') (page 16, lines 15-16)* by using reliable datagram (RD) sessions *(page 18, lines 3-5)*.

Claim 17 defines a method for transmitting data in a data storage system with at least two RAID controllers and at least two L4 routers *(See Figures 16 and 17)*, comprising:

determining functionality of the at least two L4 routers *(See Figure 17 and page 48 lines 3-10)*;

during a read operation *(See Figure 17 and page 49 line 15)*,

communicating the data from a storage device to a functional L4 router *(See Figure 17 and page 49 line 16)*;

determining at least one destination host for the data *(See Figure 17 and page 50 line 22)*,

transferring the data to the at least one destination host using L4 routing *(See Figure 17 and page 50 line 20- page 51line10)*;

and

during a write operation *(See Figure 17 and page 48 line 10)*

communicating the data from a host to a functional L4 router *(See Figure 17 and page 48 lines 10-12)*,

determining at least one destination storage device for the data *(See Figure 17 and page 48 line 20)*,

transferring the data to the destination storage device using L4 routing *(See Figure 17 and page 49 line 3)*, wherein the transferring includes,

generating an external queue pair (QP) *(168) (page 16, lines 1-3)* for establishing communication between the L4 router *(IB L4 Router)* and an external host *(102, 104)* through a reliable connection (RC) session *(lines from 168 to hosts 102 and 104) (page 17, lines 10-12) (page 17, lines 19-21)*, and

generating internal queue pairs *(164, pairs connected to EEC boxes) (page 17, lines 21-22) (page 18, lines 1-2)*, one of the internal queue pairs being coupled with the external queue pair*(168) (page 17, lines 10-12; page 18, lines 1-2)*, the internal queue pairs establishing communications between the at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802)* and an internal target device *(112', 114') (page 16, lines 15-16)*, between the L4 router circuitry *(IB L4 Router)* and the at least two controllers, and between the L4 router circuitry *(IB L4 Router)* and the internal target device *(112', 114') (page 16, lines 15-16)* by using reliable datagram (RD) sessions *(page 18, lines 3-5)*.

Claim 21 defines A storage network architecture *(See Figure 16)*, comprising:

at least two target devices *(Figure 16 721a-p);*

at least two controllers *(See Figure 16 controllers 108a and 108b)* for managing the at least two target devices *(Figure 16 721a-p)*, each of the at least two controllers *(See Figure 16 controllers 108a and 108b)* configured to be capable of managing the at least two target devices *(Figure 16 721a-p)* when one of the at least two controllers *(See Figure*

16 controllers 108a and 108b) is inoperable *(See Figures 16 and 17 and pages 47 lines 19-24)*;

at least two switches connecting the at least two controllers and the at least two target devices *(See Figure 16 switches 742a and 742b and controllers 108a and 108b)*; and

at least two L4 routers, *(IB L4 Router- See Figure 16 routers 162a and 162b)* each one of the at least two L4 routers being capable of communicating data between a host *(102, 104 in Figures 2-5)* and the at least two target devices *(Figure 16 721a-p)* through one of the at least two switches and one of the at least two controllers *(See Figure 16 switches 742a and 742b and controllers 108a and 108b)*, the L4 router being capable of facilitating remote direct memory access (RDMA) communications *(See Figure 1 and page 3)* between the at least two target devices *(Figure 16 721a-p)* and the host *(102, 104 in Figures 2-5)* wherein the router uses information at a transport layer to route data between transport sessions the at least_two L4 routers (*See Figure 16 routers 162a and 162b)* including,

circuitry for generating an external queue pair (QP) *(168) (page 16, lines 1-3)* for establishing communication between the L4 routers (*See Figure 16 routers 162a and 162b)* and the host *(102, 104 in Figures 2-5)*through a reliable connection (RC) session, *(lines from 168 to hosts 102 and 104) (page 17, lines 10-12) (page 17, lines 19-21)* and

circuitry for generating internal queue pairs, the internal queue pairs *(164, pairs connected to EEC boxes) (page 17, lines 21-22) (page 18, lines 1-2)* establishing communication between the L4 routers and the at least two controllers (*Figure 16 108a and 108b, Figure 17 operations 740 and 802)*, between the at least two controllers and the at least two target devices *(See Figure 16 switches 742a and 742b and controllers 108a*

and 108b), and between the L4 routers and the at least two target devices *(Figure 16 721a-p)* by using reliable datagram (RD) sessions *(page 18, lines 3-5)*.

It should be appreciated that the above description represents only a summary of the present invention. A more in-depth discussion of the present invention is provided in the Detailed Description section of the application.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Claims 1-8, 10-17, 19-22, and 30-34 have been rejected under 35 U.S.C. 103(a) as being unpatentable over Dellacona et al. (US Patent No. 6,799,224), in view of Coates et al (US Patent Application Publication No. 2005/0246393) and further in view of Neal et al. (U.S. Patent No. 6,990,528).

VII. ARGUMENT

A. Rejection under 35 U.S.C. 103(a) over Dellacona et al. in view of Coates and further in view of Neal et al.

1. <u>Claims 1-8, 10-17, 19-22, and 30-34</u>

i. Neal does not suggest generating an external queue pair (QP) for establishing communication between the L4 router and an external host through a reliable connection (RC) session, and generating internal queue pairs, one of the internal queue pairs being coupled with the external queue pair, the internal queue pairs establishing communications between the at least two controllers and an

internal target device, between the L4 router circuitry and the at least two controllers, and between the L4 router circuitry and the internal target device by using reliable datagram (RD) sessions.

The Examiner simplifies this feature in the office action as "using an external queue pair and internal queue pair to establish communication between two controllers and an internal target device through a reliable connection session." The Examiner provides no basis whatsoever for modifying the claim language for purposes of examination. The Examiner then states that Neal teaches the generating of a reliable datagram queue pair and then store the same reliable datagram domain within an end to end context within a local host channel adapter and refers to column 3, lines 17-29. As stated by the Applicants, and as continually ignored by the Examiner, column 3, lines 17-29 contain the brief description for Figures 5-10. The Examiner cites the brief description of the drawings of the reference as disclosing these features and provides no analysis whatsoever, only conclusory statements. Neal describes host consumer processes talking to devices using RD, and nowhere does Neal suggest a host using RC communications with external queue pairs, and an RD connection for internal queue pairs and an internal device, as claimed by Applicant. Neal specifies that the invention describes how SAN RD QPs can be associated with one or more EEC at the local HCA (See column 10, lines 19-24 and 58-60). Applicants would like for the Examiner to explain what this has to do with the claim language in question as the Examiner has been unable to provide a coherent response, let alone any response with rational underpinnings. Furthermore, there is no suggestion in Neal, nor any of the other cited references of mixing RC and RD communications and directing them at the router, as claimed. Neal is strictly limited to RD connections (see column 10, lines 33-43).

Dellacona, Coates, nor Neal teach having RC connections to the host and RD connections to the internal device, as claimed by Applicant. In response to the Applicant's request to supply an analysis with rationale underpinnings rather than conclusory statements based simply on word searches, the Examiner asserts that the mention of the RC messaging mechanism can simply be substituted as it is well known. The entire rejection is still based on word searches and is simply illogical.

Neal teaches that the SAN "consumers ... are processes executing on host processor node 200" (Col. 5, lines 65-66), and that "[t]he four SAN Service Types which provide this level of access control are: Reliable Connected, Unreliable Connected, Unreliable Datagram, and Raw Datagram." Neal describes the host consumer processes talking to devices using RD, thus it does not suggest a host using RC communications, as claimed by Applicant. As a matter of fact, Neal states that the RC QPs is logarithmic (P^2xN), while the scaling for RD QPs is linear (P+N). The Examiner completely ignores this fact and states that better load balancing would result from the use of the RC messaging scheme. This assertion is incorrect and completely void of any reality.

Applicants again respectfully request that the Examiner specify where combining the RC and RD connections are disclosed, as claimed in claims 1, 11, 17 and 21, in Neal or any of the other references, rather than a citation to the brief description of the drawings without providing any analysis. Applicants again further request that the Examiner provide articulated reasoning to support the conclusory statement that Neal discloses the above mentioned feature so that the Applicants may refute this reasoning on Appeal, if this rejection is maintained.

In summary, the Applicants have repeatedly asked the Examiner to provide an analysis of where in any of the references an external queue pair (QP) establishes

communication between the L4 router and an external host through a reliable connection (RC) session, and where internal queue pairs, one of the internal queue pairs being coupled with the external queue pair, the internal queue pairs establishing communications between the at least two controllers and an internal target device, between the L4 router circuitry and the at least two controllers, and between the L4 router circuitry and the internal target device by using reliable datagram (RD) sessions are taught. In response, the Applicants receive a citation to the Brief Description of the Drawings section of Neal that is simply unrelated to the claim language. Furthermore, the claim language as pending has not been examined, as the Examiner has taken it upon himself to summarize the language in a manner convenient for him to submit a rejection and without any statutory basis for not considering the language. Unfortunately, this has all been at the expense of the Applicants.

Furthermore, there is no suggestion in Dellacona, Coates, nor Neal of mixing RC and RD communications and bridging them at the router, as claimed, therefore the combination of Dellacona and Neal does not suggest all the Applicant's claims and the Office's rejection is improper.

B. Conclusion

In view of the foregoing reasons, the Applicants submit that each of the claims 1-22 and 30-34 are patentable. Therefore, the Applicants respectfully request that the Board of Patent Appeals and Interferences reverse the Examiner's rejections of the claims on appeal.

Respectfully submitted,
MARTINE PENILLA & GENCARELLA, LLP

Michael L. Gencarella
Reg. No. 44,703

710 Lakeway Drive, Suite 200
Sunnyvale, CA 94085
Telephone: (408) 749-6900
Facsimile: (408) 749-6901
Customer Number 25920

VIII. CLAIMS APPENDIX

1. A method for storing data, comprising:

transmitting a storage operation request to one of at least two controllers, the at least two controllers being capable of managing communication with a plurality of targets;

directing the storage operation request to an operational one of at least two controllers when the one of at least two controllers is inoperable through an L4 router, the directing including,

generating an external queue pair (QP) for establishing communication between the L4 router and an external host through a reliable connection (RC) session, and

generating internal queue pairs, one of the internal queue pairs being coupled with the external queue pair, the internal queue pairs establishing communications between the at least two controllers and an internal target device, between the L4 router circuitry and the at least two controllers, and between the L4 router circuitry and the internal target device by using reliable datagram (RD) sessions; and

processing the storage operation request with the operational one of the at least two controllers.

2. A method for storing data as recited in claim 1, further comprising:

communicating data for storage operations using to the at least two controllers through a network fabric interconnect.

3. A method for storing data as recited in claim 1, further comprising:

communicating data between at least one storage device and at least one host using the operational one of the at least two controllers.

4. A method for storing data as recited in claim 1, wherein the storage operation request is one of a request to store of data to a target and a request to read data from a target device, the target device being a storage device.

5. A method for storing data as recited in claim 1, wherein the directing of the storage operation request includes the storage operation request through an L4 router.

6. A method for storing data as recited in claim 5, wherein the transmitting the storage request includes communicating data between a host and at least two storage devices through one of at least two L4 routers and one of the at least two controllers wherein an operable one of the at least two L4 routers is used when one of the at least two L4 routers is inoperable, the at least two L4 routers being capable of facilitating remote direct memory access (RDMA) communications between the at least two target devices and the host wherein the router uses information at a transport layer to route data between transport sessions.

7. A method for storing data as recited in claim 2, wherein the network fabric interconnect is an InfiniBand-type fabric, and the at least two controllers are RAID controllers.

8. A method for storing data as recited in claim 7, wherein the transmitting the storage operation request includes transmitting the request through an operational one of at least two bridge chips.

9. A method for storing data as recited in claim 8, wherein the at least two bridge chips are SATA-IB chips.

10. A method for storing data as recited in claim 2, wherein the directing the storage operation request includes determining a correct path through the network fabric interconnect to a proper storage device.

11. A method for storing data, comprising:

providing a plurality of storage containers, each one of the plurality storage containers having a plurality of storage devices;

generating a plurality of storage volumes where each one of the storage volumes includes at least one storage device from each one of the plurality storage containers;

managing each of the plurality of volumes with a corresponding storage device controller; and

when the corresponding storage device controller is inoperable,

accessing data on the plurality of storage volumes through at least one operable storage device controller that is configured to access the volume managed by the inoperable storage device controller, the accessing through an L4 router, and including,

generating an external queue pair (QP) for establishing communication between the L4 router and an external host through a reliable connection (RC) session, and

generating internal queue pairs, one of the internal queue pairs being coupled with the external queue pair, the internal queue pairs establishing communications between the at least two controllers and an internal target device, between the L4 router circuitry and the at least two controllers, and between the L4 router circuitry and the internal target device by using reliable datagram (RD) sessions.

12. A method for storing data as recited in claim 11, further comprising:

assigned control of each of the plurality of storage volumes to a corresponding separate controller.

13. A method for storing data as recited in claim 11, further comprising:

when a controller for a particular storage container fails,

accessing data located on other storage containers through XOR operations.

14. A method for storing data as recited in claim 11, further comprising:

managing the plurality of storage volumes by spreading the processing of input/output requests among all of a plurality of controllers, each of the plurality of controllers controlling a corresponding one of the plurality of storage volumes.

15. A method for storing data as recited in claim 11, wherein the storage containers are RAID devices.

16. A method for storing data as recited in claim 11, wherein the storage devices are disk drives.

17. A method for transmitting data in a data storage system with at least two RAID controllers and at least two L4 routers, comprising:

determining functionality of the at least two L4 routers;

during a read operation,

communicating the data from a storage device to a functional L4 router;

determining at least one destination host for the data,

transferring the data to the at least one destination host using L4 routing; and

during a write operation

communicating the data from a host to a functional L4 router,

determining at least one destination storage device for the data,

transferring the data to the destination storage device using L4 routing, wherein the transferring includes,

generating an external queue pair (QP) for establishing communication between the L4 router and an external host through a reliable connection (RC) session, and

generating internal queue pairs, one of the internal queue pairs being coupled with the external queue pair, the internal queue pairs establishing communications between the at least two controllers and an internal target device, between the L4 router circuitry and the at least two controllers, and between the L4 router circuitry and the internal target device by using reliable datagram (RD) sessions.

18. A method for transmitting data in a storage system as recited in claim 17, wherein communicating the data from a storage device to a functional L4 router and communicating the data from the host to the functional L4 router includes transmitting the data through an SATA-IB bridge and a network fabric.

19. A method for transmitting data in a data storage system as recited in claim 17, wherein the network fabric is an InfiniBand-type fabric.

20. A method for transmitting data in a data storage system as recited in claim 17, wherein the L4 routing includes using RDMA data transfer.

21. A storage network architecture, comprising:

at least two target devices;

at least two controllers for managing the at least two target devices, each of the at lest two controllers configured to be capable of managing the at least two target devices when one of the at least two controllers is inoperable;

at least two switches connecting the at least two controllers and the at least two target devices; and

at least two L4 routers, each one of the at least two L4 routers being capable of communicating data between a host and the at least two target devices through one of the at least two switches and one of the at least two controllers, the L4 router being capable of facilitating remote direct memory access (RDMA) communications between the at least two target devices and the host wherein the router uses information at a transport layer to route data between transport sessions the at least_two L4 routers including,

circuitry for generating an external queue pair (QP) for establishing communication between the L4 routers and the host through a reliable connection (RC) session, and

circuitry for generating internal queue pairs, the internal queue pairs establishing communication between the L4 routers and the at least two controllers, between the at least two controllers and the at least two target devices, and between the L4 routers and the at least two target devices by using reliable datagram (RD) sessions.

22. A network architecture as recited in claim 21, wherein each one of the at least two L4 routers utilizes at least one internal queue pair to communicate with at least one external queue pair of the host.

Claims 23-29 cancelled

30. A network architecture as recited in claim 21, wherein the at least two controllers are RAID controllers.

31. A network architecture as recited in claim 21, wherein the plurality of targets is a plurality of storage devices.

32. A network architecture as recited in claim 21, wherein the plurality of storage devices is a plurality of disk drives.

33. A network architecture as recited in claim 21, wherein when one of the at least two L4 routers is inoperable, the inoperable one of the least two L4 routers is capable of communicating data between the host and at least two target devices.

34. A network architecture as recited in claim 21; wherein one of the at least two switches is inoperable, the inoperable one of the at least two switches is capable of communicating data between the host and at least two target devices.

IX. EVIDENCE APPENDIX

There is currently no evidence entered and relied upon in this Appeal.

X. RELATED PROCEEDINGS APPENDIX

There are currently no decisions rendered by a court or the Board in any proceeding identified in the Related Appeals and Interferences section.